

MIT SLOAN SCHOOL OF MANAGEMENT

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY (CSAIL)

ARTIFICIAL INTELLIGENCE: IMPLICATIONS FOR BUSINESS STRATEGY

ONLINE SHORT COURSE

MODULE 3 UNIT 1
Video 5 Transcript

Module 3 Unit 1 Video 5 Transcript

THOMASA MALONE (TM): Another question is, we've talked mostly about understanding what's in a document or in a natural language corpus whether that's spoken or written. We also need to often generate natural language. Sometimes you need to just train it the other way, sometimes I think it's also possible to use some kind of template, right? Can you tell us a little bit about that?

REGINA BARZILAY: So, there are two ways to do generation. One way is to do generation completely from scratch. Like in machine translation, you've got the sentence in English and you generate a sentence in French. You need to select all the words, you know, string them together in the right order to get yourself a grammatical French sentence. And for this again you're training it on huge, huge amounts of data. And the model train may be on hundreds of millions of French sentences, kind of learns what constitutes a valid, grammatically correct French sentence. So, this is actually a tough task, the generation. Those of you who use the machine translation systems know that the outputs sometimes are corrupted and there are mistakes. Another approach, which is easier typically for machines to do, is to do template generation. It's like, you know, people who are weak at writing typically kind of use templated writing. So, the idea is that, let's say you want to reply to the customer and there are several possible ways to reply. Then you just need to fill in the slots. So, machine needs to first classify what kind of template to use as a response, and what information to fill in the slot. And those are two classification tasks. If it is provided with a large amount of data, it can learn it and then we are guaranteed that the output actually looks grammatically correct. And one way where people use it quite successfully is actually in financial markets, where people need to produce some kind of standard recording, based on their regulations. There is not much, you know, creative writing there. You just need to fill in the template. And now it is automated, and machines can do that.

TM: Great. So, that's an example of using a template to generate the output. And I think what you've said, in multiple ways so far, is that if you have enough examples of the inputs and the outputs you want, the system can learn that using machine learning techniques. Let's now push to something that systems can't actually do today and that is to use kind of deeper, more generalized kinds of reasoning. For instance, an example we might say is if you wanted to build a system that was kind of like a butler, and that could make useful suggestions for you whatever you're doing. And you said something like, I have a headache, and you'd like your system to be able to say, "She has a headache. Maybe she wants some medicine. She doesn't have any in the cabinet. She's going to need to buy some. That means she has to go a store and the stores are only open until six." Say the store nearby, a drugstore nearby. So, for a system to do all that kind of reasoning, I think you would say can't be done today unless you've got lots and lots of examples of that.

RB: So, there are actually systems that kind of doing this kind of learning, that they observe you and learn from your actions. I would say that if you are, you know, dealing with a new person that you haven't seen you wouldn't know when they say they have a headache that

you actually need to go to a drugstore. Maybe the person never takes drugs. So, the only reason you know, because you were many times in this situation and that these people that you observed typically go and try to get themselves a drug. So, if we can imagine some system which watches and observes you, what are you doing, it may be able to make these inferences because remember the machine have an infinite capacity to remember and to generalize to some extent. So, if it is implanted in your home, and it observes what you were doing when you, let's say somebody in your household that you have a headache, and then somebody went and got you a drug, it may be eventually able to make this kind of inferences. And we are seeing it not exactly there, but Alexa, Google Home, and other devices, they still don't do that, but they do a lot of things, and people can add new skills to them. We're just at the beginning of the road.

TM: To make this example more concrete, suppose you said, "Alexa, whenever I have a headache I want to buy medicine." Alexa wouldn't be able to figure that out, and know what to do in the future, right?

RB: Right now, you need to program a skill to Alexa, which does it, but this we are talking about the current Alexa, which is, you know, the very early generation of Alexas. There is already a lot of research work where actually there is this apprentice kind of training where the machine observes, records and sees what are you doing, and then it makes you suggestions. So, I think as we move forward, this will be the desired way of training the system. Instead of sitting and annotating millions of examples, you want the machine to nonintrusively sit and observe what human is doing and learn from this experience, the same way as, you know, another human would be learning.

TM: But, you'd also like to be able to let the machine understand from someone just telling it, and that's what's not easy today.

RB: Exactly true, right now it needs to have enough training examples to understand, and I think it's just a matter of our learning technology. Today it's at the point that it needs to have a very explicit supervision. Where I believe machine learning will be going is to move away – and natural language processing – is to move away from this very concrete and specific annotations to much more abstract teaching. And then the machine can understand it much better.

TM: Great. So, I think that's actually a nice summary of where we are today, which is you have to learn from a whole, whole, whole lot of examples to where we may be headed in the future, which is the machines can learn from more abstract and explicit teaching.

RB: Yeah. Excellent.

TM: Thank you very much, Regina. This was a fascinating interview.

RB: Thank you very much.

TM: Did you understand all the concepts covered in this video? If you'd like to go over any of the sections again, please click on the relevant button.